

Testable uniqueness conditions for empirical assessment of undersampling levels in total variation-regularized x-ray CT

Jakob S. Jørgensen*, Christian Kruschel†, and Dirk A. Lorenz†

September 2, 2014

Abstract

We study recoverability in fan-beam computed tomography (CT) with sparsity and total variation priors: how many underdetermined linear measurements suffice for recovering images of given sparsity? Results from compressed sensing (CS) establish such conditions for, e.g., random measurements, but not for CT. Recoverability is typically tested by checking whether a computed solution recovers the original. This approach cannot guarantee solution uniqueness and the recoverability decision therefore depends on the optimization algorithm. We propose new computational methods to test recoverability by verifying solution uniqueness conditions. Using both reconstruction and uniqueness testing we empirically study the number of CT measurements sufficient for recovery on new classes of sparse test images. We demonstrate an average-case relation between sparsity and sufficient sampling and observe a sharp phase transition as known from CS, but never established for CT. In addition to assessing recoverability more reliably, we show that uniqueness tests are often the faster option.

Keywords: Computed tomography; total variation; sparse regularization; uniqueness conditions.

1 Introduction

Regularization methods for tomographic reconstruction that exploit sparsity have been in the focus of research recently. Motivated by the theory of compressed sensing (CS) [1, 2] many papers proposed to use sparse or total variation (TV) regularization to compute tomographic reconstructions from underdetermined measurements [3, 4, 5]. One promising goal is unchanged or even improved reconstruction quality from a significantly reduced sampling effort, thereby lowering the necessary radiation dose in medical computed tomography (CT) and scanning time in e.g. materials science and non-destructive testing.

Compressed sensing offers methodologies to predict under what circumstances it is possible to compute exact reconstructions from underdetermined linear measurements. Usually these conditions depend both on the measurement matrix and on the signal class that is considered. Roughly spoken, a standard result reads as follows: All vectors that are sparse enough can be reconstructed exactly from underdetermined measurements with a random matrix (e.g. all entries independently, identically Gaussian distributed) by computing the solution of the linear system that has the smallest ℓ^1 -norm [6]. There are also results [7] that state exact recovery conditions for TV regularization [8]. An overview of CS recovery guarantees can be found in [9]. It is generally acknowledged, however, that existing guarantees either do not apply to or give extremely pessimistic bounds in deterministic sampling contexts [10]. In particular for CT, [11, 12, 13] describe the lack of general guarantees, while [11, 12] derive preliminary average-case results for certain restricted special geometries known as discrete geometry; however these results do not cover regular sampling patterns in CT, such as parallel-beam and fan-beam geometries.

In our recent work [14, 13] we have been interested in establishing conditions on sparsity and sampling levels sufficient for image recovery with regular CT sampling patterns. In particular, [14] suggests a link between gradient sparsity and sufficient sampling for accurate TV-reconstruction. In [13], we carried out

*Corresponding author. Email: jakj@dtu.dk. Department of Applied Mathematics and Computer Science, Technical University of Denmark, Richard Petersens Plads, 2800 Kgs. Lyngby, Denmark.

†Institute for Analysis and Algebra, Technical University of Braunschweig, Pockelsstraße 14, 38106 Braunschweig, Germany.

empirical studies of the average sufficient number of CT fan-beam projection views for ℓ^1 -recovery as function of image sparsity. Using a phase diagram similar to the Donoho-Tanner [15] phase diagram we showed that ℓ^1 -recovery often admits sharp phase transitions as the sampling level is increased, and that the critical sampling level increases with the number of image nonzeros.

The present work considerably expands on the results of [13] by addressing two limitations. First, while ℓ^1 -regularization is useful for CT, TV-regularization is often a more successful sparsity prior for CT, because many objects have a piecewise constant appearance. The present work extends to study recovery using anisotropic and isotropic TV-regularization and as a step towards this proposes ways to generate images of a desired gradient sparsity. Second, both ℓ^1 and TV reconstruction is possibly subject to non-unique solutions. The approach of [13] considers an image to be uniquely recoverable if numerical solution of the relevant optimization problem recovers the original image, but does not consider whether the computed solution is unique. In the present work, we derive uniqueness tests that can be used to computationally verify ℓ^1 and TV solution uniqueness. We then compute phase diagrams using both reconstruction and uniqueness tests to verify the ℓ^1 recovery results from [13] and for anisotropic and isotropic versions of TV. In all cases, we observe a pronounced average-case relation between sparsity and the sufficient sampling level for recovery as well as a sharp phase transition. We also compare the reconstruction and uniqueness test approaches in terms of computing time.

The paper is organized as follows. Section 2 describes the CT imaging model and ℓ^1 -norm and total variation regularization problems in study. Section 3 describes necessary and sufficient conditions for solution uniqueness, while Section 4 presents our numerical implementations of reconstruction and uniqueness tests. Section 5 describes how to generate test images with desired image or gradient sparsity. Section 6 presents our results establishing empirically a relation between sparsity and the average sufficient sampling level for recovery. Finally, Section 7 discusses the results and concludes the paper.

2 Sparse image reconstruction methods for computed tomography

2.1 Imaging model

Imaging by computed tomography (CT) exploits that x-rays are attenuated when passing through matter. The attenuation depends on the material traversed by the x-ray, as described by the so-called linear attenuation coefficient μ , with denser materials generally attenuating more. The intensities I_0 and I of an x-ray before and after passing through an object with linear attenuation coefficient $\mu(\mathbf{s})$, as function of the spatial coordinate \mathbf{s} , can be modeled by Lambert-Beer's law, see, e.g., [16], which in a rearranged form reads

$$-\log \frac{I}{I_0} = \int_L \mu(\mathbf{s}) d\mathbf{s}, \quad (1)$$

where $\int_L \mu(\mathbf{s}) d\mathbf{s}$ denotes the line integral along the line L describing the x-ray path. By means of discretizing the object into n pixels and the data by assuming that m individual x-rays with infinitesimal width are used, a fully discretized imaging model can be written $\mathbf{b} = \mathbf{A}\mathbf{x}$. Here, \mathbf{x} is a vector of length n of all pixel values stacked. \mathbf{A} is an m -by- n matrix of which the (i, j) th element, \mathbf{A}_{ij} , equals the path length of the i th ray through pixel j , such that $\sum_j \mathbf{A}_{ij} \mathbf{x}_j$ approximates the line integral in (1) for ray i . Each ray only intersects a small number of pixels (for a square \sqrt{n} -by- \sqrt{n} array of the order of \sqrt{n}), causing the remaining \mathbf{A}_{ij} values to be zero and \mathbf{A} to be very sparse. \mathbf{b} is a vector of length m with the log-transformed data, i.e., $\mathbf{b}_i = -\log(I_i/I_0)$ for rays $i = 1, \dots, m$.

In the present work, we consider a 2-D fan-beam geometry with equi-angular projection views acquired from 360° around the image. Due to rotational symmetry we consider the image to be the largest inscribed disk within a square N_{side} -by- N_{side} pixel array, hence consisting of $n \approx \pi/4 \cdot N_{\text{side}}^2$ pixels. The source-to-detector distance is set to $2N_{\text{side}}$ and the detector has the shape of a circular arc centered at the source and consists of $2N_{\text{side}}$ detector elements. The number of projection views is denoted N_v and the fan angle is set to 28.07° so that precisely the inscribed disk is covered. The total number of linear measurements is $m = 2N_{\text{side}}N_v$ and the m -by- n system matrix \mathbf{A} is computed using the function `fanbeamtomo` in the MATLAB package AIR Tools [17].

2.2 Sparse regularization

In the context of an underdetermined system $\mathbf{b} = \mathbf{A}\mathbf{x}$, $\mathbf{A} \in \mathbb{R}^{m \times n}$, $m < n$, one obtains a whole affine space of solutions. One selects one solution of this space by considering a regularization functional R , and computing the minimum- R solution, i.e. the solution of

$$\min_{\mathbf{x}} R(\mathbf{x}) \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{b}.$$

In the present work we study ℓ^1 -norm (L1), anisotropic TV (ATV), and isotropic TV (ITV) regularization. First, the ℓ^1 -norm, defined as

$$\text{L1 : } \quad R(\mathbf{x}) = \|\mathbf{x}\|_1 = \sum_j |\mathbf{x}_j|, \quad (2)$$

enforces sparsity of the minimizer. ℓ^1 -norm minimization forms one backbone of compressed sensing and the problem of computing minimum- ℓ^1 -norm solution of underdetermined systems is also known as Basis Pursuit [18].

Second, anisotropic TV can be written using a set of vectors \mathbf{d}_i , $i = 1, \dots, N$ (of length n as \mathbf{x}) and the matrix $\mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_N]$ as

$$\text{ATV : } \quad R(\mathbf{x}) = \|\mathbf{D}^T \mathbf{x}\|_1 = \sum_{i=1}^N |\mathbf{d}_i^T \mathbf{x}|. \quad (3)$$

The matrix \mathbf{D} is called dictionary in this context and the minimum- R solution is seeking a solution in which inner products with the dictionaries entries form a sparse vector, for example to enforce sparsity in the coefficients of a wavelet basis or a dictionary learned from training images. We use ATV to denote the general case but focus in the present work on the anisotropic TV, which is a special case where \mathbf{D}^T contains finite-difference approximations of the horizontal and vertical derivatives in each pixel, i.e., $N = 2n$. However, we emphasize that all our theoretical results hold in the general case.

Third, isotropic TV can be written as a special case of the group sparsity problem [19]. This problem differs in a small but crucial point from ATV: Here we do not simply take the ℓ^1 -norm of $\mathbf{D}^T \mathbf{x}$ as objective function, but we take a mixed $\ell^{1,2}$ -norm. To fix notation, consider a linear mapping $\mathcal{D} : \mathbb{R}^{r \times p} \rightarrow \mathbb{R}^n$, i.e. the transposed map is $\mathcal{D}^T : \mathbb{R}^n \rightarrow \mathbb{R}^{r \times p}$, where r is the number of groups and p is the number of pixels in each group. For $\mathbf{Y} \in \mathbb{R}^{r \times p}$ we consider the mixed $\ell^{1,2}$ -norm

$$\|\mathbf{Y}\|_{1,2} = \sum_{i=1}^r \|\mathbf{Y}_i\|_2 = \sum_{i=1}^r \left(\sum_{j=1}^p |\mathbf{Y}_{i,j}|^2 \right)^{1/2},$$

where $\mathbf{Y}_i \in \mathbb{R}^p$ denotes the i -th row of \mathbf{Y} . We then set

$$\text{ITV : } \quad R(\mathbf{x}) = \|\mathcal{D}^T \mathbf{x}\|_{1,2} = \sum_{i=1}^r \|(\mathcal{D}^T \mathbf{x})_i\|_2, \quad (4)$$

We use ITV to denote the general case but focus in the present work on the isotropic TV, for which we take each group to be the horizontal and vertical finite-difference partial-derivative approximations in each pixel, i.e., $p = 2$ and $r = n$. Again, our theoretical results hold in the general case.

Note that the matrix \mathbf{D} in ATV and linear mapping \mathcal{D} in ITV are closely related for the case of finite differences, namely

$$(\mathcal{D}^T \mathbf{x})_i = [\mathbf{d}_i^T \mathbf{x}, \mathbf{d}_{i+n}^T \mathbf{x}] \quad \text{for all } i \leq n.$$

3 Necessary and sufficient conditions for uniqueness of minimizer

In this section, conditions for the uniqueness of the considered optimization problems are introduced. We use the following notation. The complement of an index set $I \subset \{1, \dots, n\}$ is denoted as $I^c = \{1, \dots, n\} \setminus I$. If $\mathbf{A} \in \mathbb{R}^{m \times n}$, $m < n$, then \mathbf{A}_I denotes the submatrix of \mathbf{A} with columns indexed by I . The transposed of such a submatrix is denoted by \mathbf{A}_I^T .

3.1 Uniqueness conditions for L1

The following theorem gives necessary and sufficient conditions for a vector \mathbf{x}^* to solve L1 uniquely.

Theorem 3.1 ([20]). *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m < n$ and $\mathbf{x}^* \in \mathbb{R}^n$ with $I = \text{supp}(\mathbf{x}^*)$. Then \mathbf{x}^* is the unique solution of*

$$\text{L1 : } \min_{\mathbf{y}} \|\mathbf{y}\|_1 \text{ subject to } \mathbf{A}\mathbf{y} = \mathbf{A}\mathbf{x}^* \quad (5)$$

if and only if

$$\ker(\mathbf{A}_I) = \{\mathbf{0}\} \quad (6)$$

and there exists $\mathbf{w} \in \mathbb{R}^m$ such that

$$\mathbf{A}_I^T \mathbf{w} = \text{sign}(\mathbf{x}^*)_I, \quad \|\mathbf{A}_{I^c}^T \mathbf{w}\|_\infty < 1. \quad (7)$$

The theorem has two important consequences: First, the recoverability of a vector \mathbf{x}^* only depends on its sign pattern, not the magnitude of its entries. Second, to know if some vector \mathbf{x}^* can be recovered from the measurement $\mathbf{A}\mathbf{x}^*$ (for some given, fixed \mathbf{A}), one only needs to check the existence of a vector $\mathbf{w} \in \mathbb{R}^m$ such that the equality and the inequality from (7) are fulfilled. As we will show in Section 5, this can be done by solving an m -dimensional linear program. Since the vector \mathbf{w} is related to the dual optimization problem of (5), it will be called *dual certificate* in the following.

A seemingly different necessary and sufficient condition for \mathbf{x}^* being the unique L1 solution given in [21] is the existence of a vector \mathbf{w} such that $\mathbf{A}_I^T \mathbf{w} = \text{sign}(\mathbf{x}^*)_I$, $\|\mathbf{A}_{I^c}^T \mathbf{w}\|_\infty \leq 1$ and that \mathbf{A}_J is injective where $J = \{j \in \{1, \dots, n\} \mid |(\mathbf{A}^T \mathbf{w})_j| = 1\}$. This condition may appear weaker but is in fact equivalent to Theorem 3.1 and we do not use it any further.

3.2 Uniqueness conditions for ATV

We extend the previous result to the ATV case through a straightforward generalization of Theorem 3.1. The following theorem basically adapts the result in [22].

Theorem 3.2. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m < n$, $\mathbf{D} \in \mathbb{R}^{n \times N}$ and $\mathbf{x}^* \in \mathbb{R}^n$ with $I = \text{supp}(\mathbf{D}^T \mathbf{x}^*)$. Then it holds that \mathbf{x}^* is the unique solution of*

$$\text{ATV : } \min_{\mathbf{y}} \|\mathbf{D}^T \mathbf{y}\|_1 \text{ subject to } \mathbf{A}\mathbf{y} = \mathbf{A}\mathbf{x}^* \quad (8)$$

if and only if

$$\ker(\mathbf{A}) \cap \ker(\mathbf{D}_{I^c}^T) = \{\mathbf{0}\} \quad (9)$$

and there exists $\mathbf{w} \in \mathbb{R}^m$ and $\mathbf{v} \in \mathbb{R}^N$ such that

$$\mathbf{D}\mathbf{v} = \mathbf{A}^T \mathbf{w}, \quad \mathbf{v}_I = \text{sign}(\mathbf{D}_I^T \mathbf{x}^*), \quad \|\mathbf{v}_{I^c}\|_\infty < 1.$$

Proof. The proof is separated into two parts, each for one direction. First, it will be shown that \mathbf{x}^* is the unique solution under the given conditions. For each $\mathbf{y} \in \mathbb{R}^n$ with $\mathbf{y} \neq \mathbf{x}^*$ it holds $\mathbf{D}_{I^c}^T \mathbf{y} = \mathbf{D}_{I^c}^T (\mathbf{y} - \mathbf{x}^*) \neq \mathbf{0}$ since $\mathbf{y} - \mathbf{x}^*$ is a nontrivial null space element of \mathbf{A} and (9) is provided. Further, with $\mathbf{s} = \text{sign}(\mathbf{D}^T \mathbf{x}^*)$, the remaining conditions imply

$$\begin{aligned} \|\mathbf{D}^T \mathbf{x}^*\|_1 &= \mathbf{s}^T \mathbf{D}^T \mathbf{x}^* = (\mathbf{D}\mathbf{v})^T \mathbf{x}^* = \mathbf{w}^T \mathbf{A}\mathbf{x}^* = \mathbf{w}^T \mathbf{A}\mathbf{y} = \mathbf{v}^T \mathbf{D}^T \mathbf{y} \\ &\leq \underbrace{\|\mathbf{v}_I\|_\infty}_{=1} \|\mathbf{D}_I^T \mathbf{y}\|_1 + \underbrace{\|\mathbf{v}_{I^c}\|_\infty}_{<1} \|\mathbf{D}_{I^c}^T \mathbf{y}\|_1 < \|\mathbf{D}^T \mathbf{y}\|_1. \end{aligned}$$

As a result of (9), the latter inequality is truly strict, see above. This proves the conditions to be sufficient for ATV uniqueness.

Let us assume the vector \mathbf{x}^* solves the considered optimization problem uniquely. Since the optimization problem (8) is piecewise linear, for all $\mathbf{h} \in \ker(\mathbf{A}) \setminus \{\mathbf{0}\}$ it holds that

$$0 < \lim_{t \rightarrow 0, t > 0} \frac{1}{t} [\|\mathbf{D}^T(\mathbf{x}^* + t\mathbf{h})\|_1 - \|\mathbf{D}^T\mathbf{x}^*\|_1] = \sum_{i \in I} \text{sign}(\mathbf{d}_i^T \mathbf{x}^*) \mathbf{d}_i^T \mathbf{h} + \sum_{i \in I^c} |\mathbf{d}_i^T \mathbf{h}|,$$

and, since the latter inequality holds for all nontrivial null space elements, in particular $-\mathbf{h}$, it holds that

$$\left| \sum_{i \in I} \text{sign}(\mathbf{d}_i^T \mathbf{x}^*) \mathbf{d}_i^T \mathbf{h} \right| < \sum_{i \in I^c} |\mathbf{d}_i^T \mathbf{h}|. \quad (10)$$

Moreover, it holds $\mathbf{D}_{I^c}^T \mathbf{h} \neq \mathbf{0}$ which implies (9).

Finally the rest of the conditions will be proved. Choose $\boldsymbol{\eta} \in \mathbb{R}^N$ with $\eta_i = \text{sign}(\mathbf{d}_i^T \mathbf{x}^*)$ for $i \in I$ and $\eta_j = 0$ for $j \notin I$ and assume $\mathbf{D}^T \boldsymbol{\eta}$ is not an element of the range of \mathbf{A}^T – otherwise the proof would be finished. Let $\ker \mathbf{A}$ be p -dimensional and $\{\mathbf{w}^{(l)}\}_{1 \leq l \leq p}$ be a basis of $\ker \mathbf{A}$ with $1 = \boldsymbol{\eta}^T \mathbf{D}^T \mathbf{w}^{(l)}$; it holds

$$1 = \boldsymbol{\eta}^T \mathbf{D}^T \mathbf{w}^{(l)} = \sum_{i \in I} \text{sign}(\mathbf{d}_i^T \mathbf{x}^*) \mathbf{d}_i^T \mathbf{w}^{(l)} \text{ for all } 1 \leq l \leq p.$$

In the following, a vector $\tilde{\boldsymbol{\xi}} \in \mathbb{R}^N$ will be constructed such that $\tilde{\boldsymbol{\xi}}^T \mathbf{D}^T \mathbf{w}^{(l)} = -1$ holds for all $1 \leq l \leq p$; hence, the vector $\mathbf{D}^T(\boldsymbol{\eta} + \tilde{\boldsymbol{\xi}})$ is an element of the range of \mathbf{A}^T due to its orthogonality to the null space of \mathbf{A} . Consider a solution $\tilde{\boldsymbol{\xi}} \in \mathbb{R}^N$ of the problem

$$\min_{\tilde{\boldsymbol{\xi}} \in \mathbb{R}^N} \max_{j \in I^c} |\tilde{\xi}_j| \text{ subject to } \boldsymbol{\eta}^T \mathbf{D}^T \mathbf{w}^{(l)} = -1 \text{ for all } 1 \leq l \leq p$$

and $\mathbf{q}^* \in \mathbb{R}^p$ as a solution of its dual problem

$$\min_{\mathbf{q} \in \mathbb{R}^p} - \sum_{i=1}^p q_i \text{ subject to } \sum_{j \in I^c} \left| \sum_{l=1}^p q_l \mathbf{d}_j^T \mathbf{w}^{(l)} \right| \leq 1,$$

then for all $1 \leq l \leq p$ it holds

$$\max_{j \in I^c} |\tilde{\xi}_j| = \left| \sum_{i=1}^p q_i^* \right| = \left| \sum_{i \in I} \text{sign}(\mathbf{d}_i^T \mathbf{x}^*) \mathbf{d}_i^T \sum_{i=1}^p q_i^* \mathbf{w}^{(l)} \right| < \sum_{j \in I^c} \left| \sum_{l=1}^p q_l^* \mathbf{d}_j^T \mathbf{w}^{(l)} \right| \leq 1$$

since $\sum_{i=1}^p q_i^* \mathbf{w}^{(l)}$ satisfies (10). This proves the remaining conditions as necessary. \square

We also call the pair (\mathbf{v}, \mathbf{w}) of vectors a *dual certificate* for ATV. Note that adding a kernel element of \mathbf{D}^T does not affect the property of being uniquely recoverable by ATV:

Corollary 3.3. *In the setting of Theorem 3.2, let $\mathbf{h} \in \ker(\mathbf{D}^T)$ and let \mathbf{x}^* be the unique solution of (8). Then $\tilde{\mathbf{x}} = \mathbf{x}^* + \mathbf{h}$ is also a unique solution of (8), with \mathbf{x}^* replaced by $\tilde{\mathbf{x}}$.*

Proof. Just note that $\text{supp}(\mathbf{D}^T \mathbf{x}^*) = \text{supp}(\mathbf{D}^T \tilde{\mathbf{x}})$ and that if (\mathbf{v}, \mathbf{w}) is a dual certificate for \mathbf{x}^* it is also a dual certificate for $\tilde{\mathbf{x}}$. \square

3.3 Uniqueness conditions for ITV

The following theorem gives sufficient conditions on \mathbf{A} , \mathcal{D} and \mathbf{x}^* such that \mathbf{x}^* is the unique ITV solution; as far as we know, it is unknown if they are also necessary.

Theorem 3.4. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$, $m < n$, $\mathcal{D} : \mathbb{R}^{r \times p} \rightarrow \mathbb{R}^n$ and $\mathbf{x}^* \in \mathbb{R}^n$ and denote by $I = \{i \in \{1, \dots, r\} : (\mathcal{D}^T \mathbf{x}^*)_i \neq 0\}$. Then it holds that \mathbf{x}^* is the unique solution of*

$$\text{ITV : } \min_{\mathbf{y}} \|\mathcal{D}^T \mathbf{y}\|_{1,2} \text{ subject to } \mathbf{A} \mathbf{y} = \mathbf{A} \mathbf{x}^* \quad (11)$$

if there exists $\mathbf{w} \in \mathbb{R}^m$ and $\mathbf{Y} \in \mathbb{R}^{r \times p}$ such that $\mathcal{D} \mathbf{Y} = \mathbf{A}^T \mathbf{w}$ and

1. There exists $\mathbf{Y} \in \mathbb{R}^{r \times p}$ such that

$$\mathbf{Y}_i = \frac{(\mathcal{D}^T \mathbf{x}^*)_i}{\|(\mathcal{D}^T \mathbf{x}^*)_i\|_2} \quad \text{for } i \in I, \quad \|\mathbf{Y}_i\|_2 < 1 \quad \text{for } i \notin I$$

and

2. with $S = \{\mathbf{v} \in \mathbb{R}^n : (\mathcal{D}^T \mathbf{v})_i = \mathbf{0} \text{ for } i \notin I\}$ it holds that $\mathbf{A}\mathbf{v} = \mathbf{A}\mathbf{y}$ implies $\mathbf{y} = \mathbf{v}$ for all $\mathbf{y}, \mathbf{v} \in S$.

Proof. On $\mathbb{R}^{r \times p}$ we have the usual inner product $\langle \mathbf{Y}, \mathbf{Z} \rangle = \sum_{i=1}^r \sum_{j=1}^p \mathbf{Y}_{ij} \mathbf{Z}_{ij}$. Then it holds for the \mathbf{Y} defined above that

$$\|\mathcal{D}^T \mathbf{y}\|_{1,2} = \langle \mathbf{Y}, \mathcal{D}^T \mathbf{y} \rangle = \langle \mathcal{D} \mathbf{Y}, \mathbf{y} \rangle = \langle \mathbf{A}^T \mathbf{w}, \mathbf{y} \rangle = \langle \mathbf{w}, \mathbf{A} \mathbf{y} \rangle.$$

Also note that obviously, $\mathbf{y} \in S$.

Now consider a vector $\mathbf{v} \in \mathbb{R}^n$ that is feasible for (11), i.e. it holds that $\mathbf{A}\mathbf{v} = \mathbf{A}\mathbf{x}^*$. If \mathbf{v} would be an element of S , then, by assumption, $\mathbf{v} = \mathbf{x}^*$. Hence, if $\mathbf{v} \neq \mathbf{y}$, then $\mathbf{v} \notin S$, i.e. there is at least one index $i_0 \notin I$ such that $(\mathcal{D}^T \mathbf{v})_{i_0} \neq \mathbf{0}$. Consequently we get

$$\begin{aligned} \|\mathcal{D}^T \mathbf{x}^*\|_{1,2} &= \langle \mathbf{w}, \mathbf{A} \mathbf{y} \rangle = \langle \mathbf{w}, \mathbf{A} \mathbf{v} \rangle = \langle \mathbf{A}^T \mathbf{w}, \mathbf{v} \rangle = \langle \mathcal{D} \mathbf{Y}, \mathbf{v} \rangle \\ &= \sum_{i=1}^r \mathbf{Y}_i \cdot (\mathcal{D}^T \mathbf{v})_i \quad [\cdot \text{ denotes the inner product in } \mathbb{R}^d] \\ &\leq \sum_{i=1}^r \|\mathbf{Y}_i\|_2 \|(\mathcal{D}^T \mathbf{v})_i\|_2 = \underbrace{\|\mathbf{Y}_{i_0}\|_2}_{<1} \underbrace{\|(\mathcal{D}^T \mathbf{v})_{i_0}\|_2}_{>0} + \sum_{i \neq i_0} \underbrace{\|\mathbf{Y}_i\|_2}_{\leq 1} \|(\mathcal{D}^T \mathbf{v})_i\|_2 \\ &< \sum_{i=1}^r \|(\mathcal{D}^T \mathbf{v})_i\|_2 = \|\mathcal{D}^T \mathbf{v}\|_{1,2}. \end{aligned}$$

In other words, every feasible \mathbf{v} different from \mathbf{x}^* has a larger objective value. \square

Whereas geometrical interpretations of the conditions in Theorem 3.2 for ATV and Theorem 3.4 for ITV are difficult to establish and may require a high level of comprehension, the L1 conditions in Theorem 3.1 can be explained easily. For, say, a full rank matrix $A \in \mathbb{R}^{m \times n}$ with $m < n$ and $\mathbf{x}^* \in \mathbb{R}^n$ with $I = \text{supp}(\mathbf{x}^*)$ and $|I| = k \leq m$, the conditions (7) can be seen as the intersection of the affine space $\text{rg}(A^T)$ with the n -dimensional hypercube $[-1, +1]^n$. Indeed, the affine space cuts the interior of an $(n - k)$ -dimensional face of the hypercube; which face is sliced depends on the sign pattern of \mathbf{x}^* .

4 Numerical implementation of reconstruction and uniqueness tests

4.1 Reconstruction problems

The three regularized reconstruction problems L1, ATV and ITV are solved numerically by a primal-dual interior-point method using MOSEK [23]. Our motivation for this choice of method is to ensure that the optimization problem is solved accurately. MOSEK achieves this by producing a certificate of optimality for the returned numerical solution, i.e., a primal-dual solution pair with duality gap numerically close to zero.

Across the large number of optimizations done for the present study, it is our experience that other methods such as accelerated first-order methods [24, 25] and primal-dual methods [26, 27] are less reliable in actually arriving at a solution of the equality-constrained problem accurate enough to reliably assess whether it is equal to the original image.

We solve L1 as a linear program (LP) by introducing $\mathbf{q} \in \mathbb{R}^n$ to bound \mathbf{x} :

$$\min_{\mathbf{x}, \mathbf{q}} \mathbf{1}^T \mathbf{q} \quad \text{subject to} \quad \mathbf{A} \mathbf{x} = \mathbf{b} \quad \text{and} \quad -\mathbf{q} \leq \mathbf{x} \leq \mathbf{q}. \quad (12)$$

In a similar fashion, ATV can be solved as an LP. By defining $\mathbf{z} = \mathbf{D}^T \mathbf{x}$ and using \mathbf{q} for bounding \mathbf{z} we can solve the problem as

$$\min_{\mathbf{x}, \mathbf{z}, \mathbf{q}} \mathbf{1}^T \mathbf{q} \quad \text{subject to} \quad \mathbf{A} \mathbf{x} = \mathbf{b} \quad \text{and} \quad \mathbf{z} = \mathbf{D}^T \mathbf{x} \quad \text{and} \quad -\mathbf{q} \leq \mathbf{z} \leq \mathbf{q}. \quad (13)$$

ITV can be recast as a conic optimization problem, which can also be solved by MOSEK. Again, we introduce the bounding vector $\mathbf{q} \in \mathbb{R}^n$ and can solve the problem as

$$\min_{\mathbf{x}, \mathbf{q}} \mathbf{1}^T \mathbf{q} \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b} \quad \text{and} \quad \|(\mathcal{D}^T \mathbf{x})_j\|_2 \leq \mathbf{q}_j \quad \text{for } j = 1, \dots, n, \quad (14)$$

in which each of the n inequalities specify a quadratic conic constraint.

4.2 Uniqueness tests

As stated by Theorem 3.1, we can show that \mathbf{x}^* is the unique L1 minimizer if and only if \mathbf{A}_I is injective and there exists a $\mathbf{w} \in \mathbb{R}^m$ such that $\mathbf{A}_I^T \mathbf{w} = \text{sign}(\mathbf{x}^*)_I$ and $\|\mathbf{A}_{I^c}^T \mathbf{w}\|_\infty < 1$. Injectivity is tested by evaluating whether \mathbf{A}_I has full column rank. The second condition can be tested by minimizing $\|\mathbf{A}_{I^c}^T \mathbf{w}\|_\infty$ with respect to \mathbf{w} while respecting the equality constraint $\mathbf{A}_I^T \mathbf{w} = \text{sign}(\mathbf{x}^*)_I$. By splitting the infinity-norm into a two-sided inequality constraint, this problem can be solved as an LP,

$$\min_{t, \mathbf{w}} t \quad (15)$$

$$\text{subject to} \quad -t\mathbf{1} \leq \mathbf{A}_{I^c}^T \mathbf{w} \leq t\mathbf{1} \quad (16)$$

$$\mathbf{A}_I^T \mathbf{w} = \text{sign}(\mathbf{x}^*)_I, \quad (17)$$

which we accomplish by use of MOSEK's primal-dual interior-point method. For the optimal solution (t^*, \mathbf{w}^*) , by definition we have the smallest possible $t = \|\mathbf{A}_{I^c}^T \mathbf{w}\|_\infty$. If t^* is not smaller than one, then no \mathbf{w} exists with smaller t . We therefore declare \mathbf{x}^* the unique minimizer if $t < 1$, and if $t \geq 1$, \mathbf{x}^* cannot be the unique minimizer. Numerically, we test whether $t^* < 1 - \epsilon$ for $\epsilon = 10^{-5}$ to ensure that the inequality is satisfied strictly. Technically, by doing this we risk falsely rejecting solution uniqueness of some \mathbf{x}^* , namely the ones for which $1 - \epsilon < t^* < 1$. However, we found the decision on uniqueness to be robust to other choices of ϵ , so we believe this is not a problem in practice.

Regarding ATV, we can show solution uniqueness of a vector \mathbf{x}^* by use of Theorem 3.2. The condition of zero-intersection of $\ker(\mathbf{A})$ and $\ker(\mathbf{D}_{I^c}^T)$ can be checked numerically by evaluating whether the matrix $(\mathbf{A}; \mathbf{D}_{I^c}^T)$ has full rank, where semicolon means vertical concatenation. Similar to the L1 case, the second condition can be tested by solving in MOSEK the LP,

$$\min_{t, \mathbf{v}, \mathbf{w}} t \quad (18)$$

$$\text{subject to} \quad -t\mathbf{1} \leq \mathbf{v}_{I^c} \leq t\mathbf{1} \quad (19)$$

$$\mathbf{A}_I^T \mathbf{w} = \mathbf{D}_I \mathbf{v}_I + \mathbf{D}_{I^c} \mathbf{v}_{I^c} \quad (20)$$

$$\mathbf{v}_I = \text{sign}(\mathbf{D}_I^T \mathbf{x}^*), \quad (21)$$

and assessing whether the optimal t^* is smaller than $1 - \epsilon$.

For isotropic TV we can show solution uniqueness by use of Theorem 3.4. We let \mathbf{Y}_I and \mathbf{Y}_{I^c} denote \mathbf{Y} restricted to rows I and I^c , and similarly for \mathcal{D}_I and \mathcal{D}_{I^c} . Given \mathbf{x}^* we construct a \mathbf{Y} and a \mathbf{w} satisfying the requirements by solving the conic program

$$\min_{t, \mathbf{Y}, \mathbf{w}} t \quad (22)$$

$$\text{subject to} \quad \mathbf{D}_{I^c} \mathbf{Y}_{I^c} - \mathbf{A}^T \mathbf{w} = -\mathbf{D}_I \mathbf{Y}_I \quad (23)$$

$$\mathbf{Y}_i = \frac{(\mathcal{D}^T \mathbf{x}^*)_i}{\|(\mathcal{D}^T \mathbf{x}^*)_i\|_2} \quad \text{for } i \in I \quad (24)$$

$$\|\mathbf{Y}_i\|_2 \leq t \quad \text{for } i \in I^c. \quad (25)$$

If the optimal value t^* is greater than $1 - \epsilon$, then part (1) of Theorem 3.4 is not fulfilled and we cannot show uniqueness. If instead $t^* < 1 - \epsilon$ we proceed to part (2).

The set S can be equivalently described as the kernel of $\mathbf{D}_{I^c}^T$. Letting \mathbf{K} denote a matrix of basis vectors of the kernel, any $\mathbf{v} \in S$ can be represented using a coefficient vector \mathbf{c} such that $\mathbf{v} = \mathbf{K}\mathbf{c}$. We numerically test the injectivity requirement of \mathbf{A} on S by evaluating whether $\mathbf{A}\mathbf{K}$ has full column rank. If true, we have shown solution uniqueness.




Figure 1: Top: Realization of images from the **spikes** class of relative sparsity $\kappa = k/n$ values 0.1, 0.3, 0.5, 0.7, 0.9, gray-scale $[0, 1]$. Bottom: The same for the **signed-spikes** class, gray-scale $[-1, 1]$.

5 Generation of sparse test images

5.1 Images for L1

The **spikes** class will be used to test for exact recovery of sparse images from tomographic measurements by L1 (2). Since we are interested in recovery in dependence on the sparsity, we follow the usual approach [28] and build test images consisting of a given number of spikes at random positions and with entries sampled from a uniform distribution on $[0, 1]$. The **spikes** images hence are non-negative. We also consider a signed version, **signed-spikes**, with the only differences that entries are sampled from the uniform distribution on $[-1, 1]$. Fig. 1 shows example images of each class at a range of relative sparsity values.

5.2 Images for ATV

For ATV we wish to generate test images having a prescribed value of $\|\mathbf{D}^T \mathbf{x}\|_0$, where $\|\cdot\|_0$ is the cardinality. Due to the operator this is a less trivial task than in the directly sparse case. For certain operators this can be accomplished using a technique from [29] but as pointed out in that work, the technique does not apply to the finite-difference operator \mathbf{D}^T . Here, we present two methods for this purpose. The first method, **truncated-uniform**, produces test images that *in expectation* achieve the target sparsity, while the second, **alternating-projection**, produces test images of precisely the target sparsity.

5.2.1 The truncated-uniform class

The **truncated-uniform** class produces images \mathbf{x} according to the following heuristic. Given a target number of nonzeros k of the length- N vector $\mathbf{D}^T \mathbf{x}$, where \mathbf{D}^T has $|B^c|$ rows that do not correspond to differences across the image boundary, and a number F , which is the number of gray values in the image, satisfying

$$k \leq |B^c| \frac{F-1}{F}, \quad (26)$$

do the following:

1. Compute $\omega = (1 - \sqrt{1 - kF/(|B^c|(F-1))})/F$.
2. Partition the interval $[0, 1]$ into F intervals I_1, \dots, I_F where intervals $\ell = 1, \dots, F-1$ have widths $\omega_\ell = \omega$ and the last one has width $\omega_F = 1 - (F-1)\omega$, i.e. the intervals are $I_\ell = [(\ell-1)\omega, \ell\omega[$ for $\ell = 1, \dots, F-1$ and $I_F = [(F-1)\omega, 1]$. Denote by f_ℓ the midpoint of the ℓ -th interval.
3. Assign for every pixel the value f_ℓ with probability ω_ℓ .




Figure 2: Top: A set of **truncated-uniform** images of relative sparsity $\kappa = k/n$ values 0.2, 0.6, 1.0, 1.4, 1.8, gray-scale $[0, 1]$. Middle: A set of ‘**alternating-projection** (anisotropic) images of relative sparsity values 0.2, 0.6, 1.0, 1.4, 1.8, gray-scale $[-1, 1]$. Bottom: A set of **alternating-projection** (isotropic) images of relative sparsity values 0.1, 0.3, 0.5, 0.7, 0.9, gray-scale $[-1, 1]$.

The following lemma shows that the outcome of this method is an image that in expectation has k nonzeros after application of \mathbf{D}^T as desired:

Lemma 5.1. *Let \mathbf{x} be generated by the four steps above and denote $\mathbf{z} = \mathbf{D}^T \mathbf{x}$. Then the expected value of the number of non-zero entries in \mathbf{z} is k .*

Proof. With the described procedure, the j th entry of \mathbf{x} , $j = 1, \dots, n$ is a scalar stochastic variable \mathbf{X}_j , and the corresponding vector stochastic variable \mathbf{X} has independent, identically distributed (i.i.d.) elements. We consider also the vector stochastic variable $\mathbf{Z} = \mathbf{D}^T \mathbf{X}$ and the vector stochastic indicator variable $\boldsymbol{\delta}$ with elements

$$\delta_i = \begin{cases} 1 & \text{if } \mathbf{Z}_i \neq 0, \\ 0 & \text{if } \mathbf{Z}_i = 0, \end{cases} \quad i = 1, \dots, N.$$

The total number of nonzeros in \mathbf{Z} is described by the scalar stochastic variable

$$\delta_{\text{total}} = \sum_{i=1}^N \delta_i,$$

of which we derive the expected value $E(\delta_{\text{total}})$. Using linearity of expectation and that

$$E(\delta_i) = 0 \cdot P(\mathbf{Z}_i = 0) + 1 \cdot P(\mathbf{Z}_i \neq 0) = P(\mathbf{Z}_i \neq 0), \quad i = 1, \dots, N,$$

we get

$$E(\delta_{\text{total}}) = \sum_{i=1}^N P(\mathbf{Z}_i \neq 0). \quad (27)$$

To compute $P(\mathbf{Z}_i \neq 0)$ we distinguish between two cases: \mathbf{Z}_i is a finite difference either 1) across the boundary or 2) in the interior of the image. Formally, we partition the indices $\{1, \dots, N\}$ into a boundary set B and a complementary interior set B^c .

In case 1), $i \in B$, the choice of boundary conditions (BCs) affects the probability in question. Assuming Neumann BCs, each finite difference across the boundary is zero, i.e., $P(\mathbf{Z}_i \neq 0) = 0$ for $i \in B$ corresponding to a zero row in \mathbf{D}^T at indices B . Equivalently, these rows can be removed from \mathbf{D}^T , leaving B empty. Assuming zero BCs instead, the probability is instead equal to 1, since \mathbf{Z}_i can only be zero, if the pixel value in question is 0, which happens with probability 0. In this paper, we do not consider zero BCs more.

In case 2), we apply the law of total probability over the set $\{f_1, \dots, f_F\}$,

$$P(\mathbf{Z}_i \neq 0) = \sum_{\ell=1}^F P(\mathbf{Z}_i \neq 0 | \mathbf{X}_{\tilde{i}} = f_\ell) P(\mathbf{X}_{\tilde{i}} = f_\ell),$$

where \tilde{i} denotes the index of the pixel at which \mathbf{Z}_i is evaluated. From step (3) we know that the probability that a given pixel value is f_ℓ equals ω_ℓ . For the conditional probability, we note that \mathbf{Z}_i is computed as the difference between $\mathbf{X}_{\tilde{i}}$ and a neighboring pixel's value. Since the pixel values are independent, and since $\mathbf{X}_{\tilde{i}}$ is given to be in the ℓ th interval, the probability of having a nonzero \mathbf{Z}_i equals the probability of the neighboring pixel's value being outside interval ℓ , which is $1 - \omega_\ell$. Hence,

$$\begin{aligned} P(\mathbf{Z}_i \neq 0) &= \sum_{\ell=1}^F (1 - \omega_\ell) \omega_\ell = \sum_{\ell=1}^{F-1} (1 - \omega) \omega + (1 - \omega_F) \omega_F \\ &= (F - 1) \omega (2 - F \omega) \quad \text{for } i \in B^c. \end{aligned}$$

Inserting the two cases into (27) yields

$$E(\delta_{\text{total}}) = |B^c| (F - 1) \omega (2 - F \omega).$$

Since we want $E(\delta_{\text{total}}) = k$, we solve this quadratic equation for ω . The smaller of the two solutions guarantees that the sum of all widths is not greater than 1 and is precisely ω from step (1). Further, for ω to be real-valued we require (26). \square

For the case $N_{\text{side}} = 64$, we get $n = 3228$ pixels within the disk-shaped mask. For ATV with Neumann BCs and keeping zero-rows of \mathbf{D}^T we have $N = 2n$. Due to convexity of the disk-shaped mask we can explicitly compute $|B^c| = 2n - 2N_{\text{side}}$. In our numerical studies we wish to study images sampled from the entire sparsity range between $\kappa = 0$ and $\kappa = 2$ with the maximal relative sparsity of 1.9. By taking $F = 40$ we can achieve images \mathbf{x} with sparsity of $k = 6169$ of $\mathbf{D}^T \mathbf{x}$. This corresponds to a relative sparsity of $\kappa = k/n = 1.911$. Examples of **truncated-uniform** images are shown in Fig. 2 for a range of relative sparsity values.

5.2.2 The alternating-projection class

Since the **truncated-uniform** class consists of images of a special structure (namely, they have a prescribed number of different gray levels) it may be that they also introduce a special behavior in the recoverability. Hence, we will consider also a different class of test images. Our goal is again to produce images \mathbf{x} such that $\|\mathbf{D}^T \mathbf{x}\|_0 = k$. We reformulate the problem as follows: Find a vector \mathbf{v} such that \mathbf{v} is in the range of \mathbf{D}^T and that $\|\mathbf{v}\|_0 = k$. If we have found such a \mathbf{v} , then we get a suitable test image \mathbf{x} by solving $\mathbf{D}^T \mathbf{x} = \mathbf{v}$. For a method to construct such a \mathbf{v} , we are inspired by the feasibility problem

$$\text{find } \mathbf{v} \in \text{rg } \mathbf{D}^T \cap \{\|\cdot\|_0 \leq k\}.$$

Although the set we are looking at is the intersection of a convex with a *non-convex* one, recent results indicate that an alternating projection approach may work [30]. Hence, we perform the following iteration:

1. Choose a random starting point $\mathbf{v}^0 \in \mathbb{R}^N$; set $j = 0$.
2. Set $\mathbf{v}^{j+\frac{1}{2}}$ as orthogonal projection of \mathbf{v}^j onto $\text{rg } \mathbf{D}^T$. With the help of the pseudoinverse $(\mathbf{D}^T)^\dagger$ of \mathbf{D}^T this is written as $\mathbf{v}^{j+\frac{1}{2}} = \mathbf{D}^T (\mathbf{D}^T)^\dagger \mathbf{v}^j$.

3. Set \mathbf{v}^{j+1} as orthogonal projection of $\mathbf{v}^{j+\frac{1}{2}}$ onto the set $\{\|\cdot\|_0 \leq k\}$: Keep the largest k entries of $\mathbf{v}^{j+\frac{1}{2}}$ and set the rest to zero. If the projection yields fewer than k nonzeros, project $\mathbf{v}^{j+\frac{1}{2}}$ on a set with higher sparsity.
4. If converged, set $\mathbf{x} = (\mathbf{D}^T)^\dagger \mathbf{v}^{j+1}$; otherwise increment j and go to step 2.

If the method converges, it is guaranteed to produce an image \mathbf{x} with the desired properties. However, in practice it does not always converge. Hence, we perform a maximum number of iterations (in the range of a few thousands) and if we do not observe convergence to a feasible point \mathbf{v} , we restart the method with a different initial point. Typically, we found that only a few restarts sufficed for producing a desired image.

We also consider a non-negative version, **alternating-projection-nonneg**, for which an image is generated from an **alternating-projection** image by shifting all pixel values by the smallest possible positive scalar such that all pixel values become non-negative.

5.3 Images for ITV

For isotropic TV we basically proceeded similarly to **alternating-projection**. However, here we considered the feasibility problem

$$\text{find } \mathbf{Y} \in \text{rg } \mathcal{D}^T \cap \{\|\cdot\|_{0,2} \leq k\},$$

where $\|\cdot\|_{0,2}$ counts the number of rows with nonzero ℓ^2 -norm. Note that the latter set are the images \mathbf{x} for which the Euclidean norm of the gradient has only k non-zero entries. Hence, we modify the iteration to:

1. Choose a random starting point $\mathbf{Y}^0 \in \mathbb{R}^{r \times p}$; set $j = 0$.
2. Set $\mathbf{Y}^{j+\frac{1}{2}}$ as orthogonal projection of \mathbf{Y}^j onto $\text{rg } \mathcal{D}^T$. With the help of the pseudoinverse $(\mathcal{D}^T)^\dagger$ of \mathcal{D}^T this is written as $\mathbf{Y}^{j+\frac{1}{2}} = \mathcal{D}^T (\mathcal{D}^T)^\dagger \mathbf{Y}^j$.
3. Set \mathbf{Y}^{j+1} as orthogonal projection of $\mathbf{Y}^{j+\frac{1}{2}}$ onto the set $\{\|\cdot\|_{0,2} \leq k\}$: Keep the k rows of $\mathbf{Y}^{j+\frac{1}{2}}$ with largest 2-norm and set the rest to zero. If the projection yields fewer than k nonzero rows, project $\mathbf{Y}^{j+\frac{1}{2}}$ on a set with higher sparsity.
4. If converged, set $\mathbf{x} = (\mathcal{D}^T)^\dagger \mathbf{Y}^{j+1}$; otherwise increment j and go to step 2.

Examples of both anisotropic and isotropic **alternating-projection** images are shown in Fig. 2 for a range of relative sparsity values. The relative sparsity values for anisotropic are chosen as twice the isotropic ones to enable a rough comparison between images from each class of ‘comparable’ sparsity.

6 Numerical experiments

As in [13] we wish to show empirically that CT image reconstruction by sparsity-exploiting methods admit sharp phase transitions as known from compressed sensing. The results in [13] only covered L1. Here, we extend to ATV and ITV and construct phase diagrams by solving the reconstruction problem as well as the uniqueness tests.

6.1 Phase diagrams for L1

In [13] it was found that reconstruction by ℓ_1 -minimization for **spikes** images yields a sharp phase transition. Here, through uniqueness testing we verify this result and further extend to the class **signed-spikes** with signed entries.

We consider images of size $N_{\text{side}} = 64$ and at each relative sparsity value $\kappa = k/n = 0.025, 0.05, 0.1, 0.2, 0.3, \dots, 0.9$ we generate 100 instances. For each instance \mathbf{x}^* , we generate synthetic CT data $\mathbf{b} = \mathbf{A}\mathbf{x}^*$ corresponding to $N_v = 1, \dots, 32$ projection views. At $N_v \leq 25$ the linear system is underdetermined, while at $N_v \geq 26$ the system matrix \mathbf{A} has full rank and \mathbf{x}^* is the unique solution no matter its sparsity. We therefore use $N_v^{\text{suf}} = 26$ as a reference point of full sampling at $N_{\text{side}} = 64$ and define the relative sampling $\mu = N_v/N_v^{\text{suf}}$. For each data set, reconstruction and uniqueness test are run. If the relative error of the




Figure 3: Phase diagrams for L1. Top row: reconstruction, bottom row: uniqueness test. Left: **spikes**, center: **signed-spikes**, right: average sufficient relative sampling point along with a 99% confidence interval at each κ -value. For both classes reconstruction and uniqueness testing agree perfectly. For **signed-spikes** the average relative sufficient sampling curve is higher than for **spikes** meaning that more projections are needed for recovery.

computed solution \mathbf{x}_{L1} w.r.t. \mathbf{x}^* is sufficiently small, i.e., $\|\mathbf{x}_{L1} - \mathbf{x}^*\|_2 / \|\mathbf{x}^*\|_2 < \epsilon$ with $\epsilon = 10^{-4}$, we declare the original perfectly recovered.

Fig. 3 shows, for both the **spikes** and the **signed-spikes** classes, reconstruction and uniqueness testing phase diagrams: each rectangle corresponds to the relative sparsity value κ at its left and the relative sampling μ at its bottom and the color indicates the fraction of instances perfectly recovered by reconstruction or deemed the unique solution by the uniqueness test. The phase diagrams are divided into a ‘full-recovery’ regime, in which all instances are uniquely recovered, and a ‘no-recovery’ regime, where all instances fail to be recovered/be unique. Further, the transition from no-recovery to full-recovery is sharp, in the sense that for all relative sparsity values adding 1–2 projection views changes the recovery rate from 0% to 100%.

For both **spikes** and **signed-spikes**, the uniqueness test phase diagrams are identical to the reconstruction phase diagrams, thereby mutually verifying correctness of each method and the attained phase diagrams.

The phase transition occurs at different sampling levels for the **spikes** and **signed-spikes** classes. This is perhaps more easily seen in the right-most plot in Fig. 3 in which the average relative sampling sufficient for recovery, i.e., the smallest value of μ at which all instances of a given sparsity are recovered, is plotted for both classes along with error bars for the 99% confidence intervals. On average **signed-spikes** require more projection views for unique recovery. This is perhaps not surprising, as having negative pixel values can lead to negative entries in the data vector \mathbf{b} , something that can not happen with non-negative pixel values due to the elements of \mathbf{A} being non-negative. Nevertheless, the phase diagram reveals quantitatively how signed entries affect recoverability.




Figure 4: Phase diagrams for ATV. Top row: reconstruction, middle row: uniqueness test. Left: **alternating-projection-nonneg**, right: **alternating-projection**. Bottom row: average relative sufficient sampling point along with a 99% confidence interval at each κ -value. No difference between classes is seen, unlike the L1 case.

6.2 Phase diagrams for ATV

In the same way as for L1 we create reconstruction and uniqueness test phase diagrams for ATV. We consider first the **alternating-projection** class as well as its non-negative version **alternating-projection-nonneg**. As the sparsity is measured after application of \mathbf{D}^T , the relative sparsity can now be in the range between 0 and 2, and in addition to the κ values in the previous section, we include now $\kappa = 1.0, 1.1, \dots, 1.9$.

The resulting reconstruction and uniqueness test phase diagrams are shown in Fig. 4. As for L1, we see a partition into ‘full-recovery’ and ‘no-recovery’ regimes separated by a sharp phase transition across 1–2 projection views. The uniqueness test phase diagrams are identical to the reconstruction phase diagrams, except for a few cases in the transition region, for example for the smallest κ values for **alternating-projection-nonneg**. We explain these minor differences by the choice of numerical threshold for assessing recovery that is chosen a priori to a constant ϵ .

Contrary to the L1 case, the phase diagrams for the signed and non-negative image class are identical, which is clearly seen in the bottom plot of Fig. 4, since the transition curves (average relative sampling for recovery) of the two classes coincide. It is interesting to observe such a fundamental difference with respect to having signed or non-negative images between L1 and ATV. We can explain the difference by the fact that the




Figure 5: Phase diagrams for ATV and **truncated-uniform**. Top left: reconstruction, bottom left: uniqueness test. Right: average relative sufficient sampling point along with a 99% confidence interval at each κ -value. Very small differences (one view more needed for recovering **truncated-uniform** images) between classes is seen.

constant vector, that is added to an **alternating-projection** image to obtain the non-negative version, is in the kernel of \mathbf{D}^T , and hence makes no difference for the ATV objective value. Further, according to Corollary 3.3, if an **alternating-projection** image is the unique ATV minimizer for its data, then the non-negative image obtained by adding a vector in the kernel of \mathbf{D}^T will be the unique minimizer for its given data.

To study whether the phase diagrams depend on the image class we repeat the experiment for the **truncated-uniform** image class. The resulting reconstruction and uniqueness testing phase diagrams are shown in Fig. 5. Again, the two phase diagrams are identical and show a sharp phase transition from the no-recovery to the full-recovery regime. As can be seen from the right-most plot of the average relative sampling for recovery for the **truncated-uniform** image class compared to the **alternating-projection** image class. The two curves are nearly identical at low and high relative sparsity values. However, in the mid-range there is a small difference corresponding to around one more projection view needed on average for the **truncated-uniform** images to be recovered.

Both **alternating-projection** and **truncated-uniform** image classes are constructed to yield images of a desired target sparsity, but in fundamentally different ways. The fact that the arising phase diagrams are so similar leaves us with the interpretation that the phase diagram and in particular the phase transition curve is governed mainly by the sparsity, while the particular image class has less influence.

6.3 Phase diagrams for ITV

We repeat the reconstruction and uniqueness test study for ITV with the **alternating-projection** image class designed for ITV. Recall that in contrast to the L1 and ATV cases, Theorem 3.4 only provides a sufficient condition of solution uniqueness. This means that, in principle, instances that are not shown to be unique solutions still might be.

For ITV, the relative sparsity with respect to image size is between 0 and 1. As for L1 we construct the phase diagram for the values $\kappa = 0.025, 0.05, 0.1, 0.2, \dots, 0.9$ and $N_v = 1, \dots, 32$, see Fig. 6. Due to the numerically more challenging conic programs of isotropic TV solution accuracy was smaller than for L1 and ATV and as a result we choose a the numerical threshold to $\epsilon = 10^{-3}$.

Once again we observe a partition into ‘full-recovery’ and ‘no-recovery’ regimes clearly separated by a

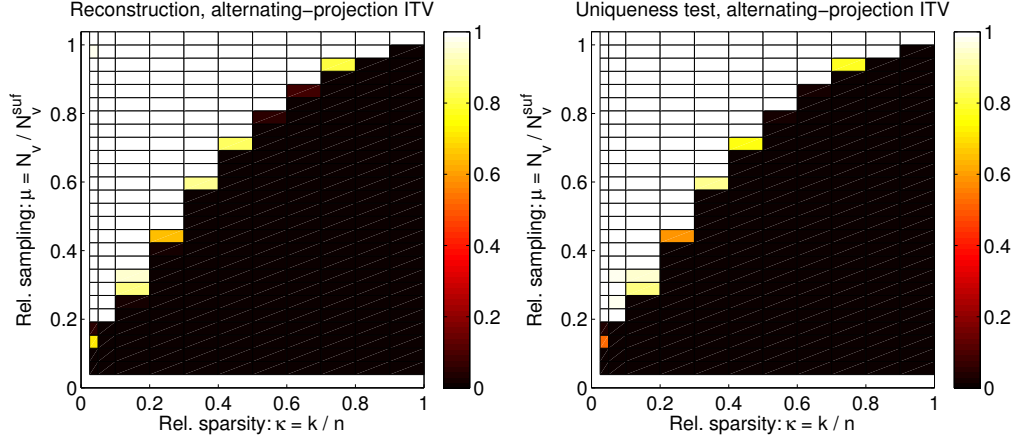


Figure 6: Phase diagrams for ITV. Left: Reconstruction, right: uniqueness test.

sharp transition. Also, the reconstruction and uniqueness test phase diagram agree almost exactly and we ascribe again the minor differences to the uniform a priori choice of the numerical threshold ϵ for assessing recovery. The almost exact agreement between the phase diagrams is interesting considering the uniqueness test is only a sufficient condition. One may conjecture that the conditions of Theorem 3.4 are in fact also necessary and that a proof is only to be found.

Even though both ATV and ITV rely on gradient sparsity, comparing their phase diagrams do not reveal a straightforward conclusion as to which method provides the greatest undersampling potential because of the different sparsity measures. More comparisons of the two methods, for example a ‘cross-over study’ of ITV reconstruction applied to the ATV image class and vice versa is beyond the scope of the present work, where the goal was to simply document phase transition behavior for CT measurements.

6.4 Computational time for reconstruction vs. uniqueness testing

We compared the computational times of reconstruction and uniqueness testing. All timing experiments were run in MATLAB 7.13 (R2011b) under Linux using MOSEK 6.0 on a Lenovo ThinkPad T430s with Intel Core i5-3320M processor (3 MB cache, up to 3.30 GHz) and 8 GB RAM, restricted to a single core.

We choose experiments with low, medium and high relative sparsity and low, medium and high relative sampling cases to measure computational times for. For the **signed-spikes** class we consider 10 instances at each of the relative sparsity levels $\kappa = 0.1, 0.3, 0.5, 0.7, 0.9$. For image size $N_{\text{side}} = 32$, we use 3, 7 and 11 views, for $N_{\text{side}} = 64$, we use 5, 13 and 21 views and for $N_{\text{side}} = 128$ we take $N_v = 9, 25$ and 41. For ATV reconstruction, we consider the **alternating-projection** image class, for the same $N_{\text{side}} = 32, 64, 128$, relative sparsity $\kappa = 0.1, 0.7, 1.3, 1.9$ and the same number of views.

Results are shown in Fig. 7. For reconstruction, computational time generally shows little dependence with κ , if any, increasing κ generally gives slightly increasing computational time. Uniqueness testing computational time tends to decrease with increasing κ . In several cases the uniqueness test is significantly faster than the reconstruction. In some of these cases, the relative sampling is low and the relative sparsity is high, which causes \mathbf{A}_I (in the L1 case) to be non-injective, and the infinity-norm minimization problem needs not be solved. In other cases, for example the L1 $N_{\text{side}} = 128$ case with $N_v = 25$, uniqueness testing is much faster across the relative sparsity range. For uniqueness testing, computational time increases with the sampling level. For reconstruction, the low sampling cases are also the fastest, however the medium sampling case is not faster than the high sampling case in all cases, for example in the L1 $N_{\text{side}} = 128$ case the times are comparable, and in the ATV $N_{\text{side}} = 32$ case, the high sampling case is in fact faster.

We conclude that in general uniqueness testing is not slower than doing reconstruction, in most cases the computational times are comparable and in some cases, uniqueness testing is in fact faster. We note that uniqueness testing can be conveniently used in case of larger κ , where reconstruction tends to be the slower option.

It is clear that the reported computational times rely on our use of MOSEK for solving the optimization

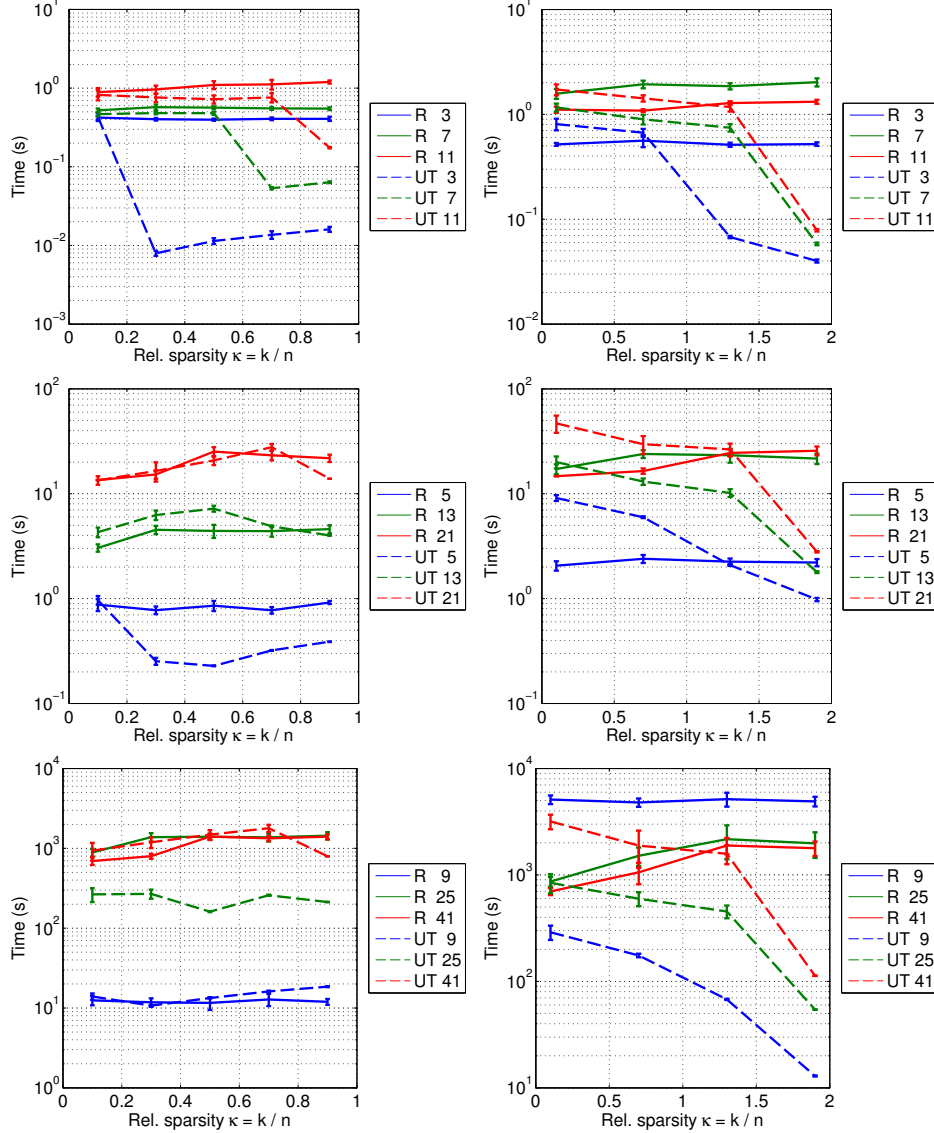


Figure 7: Average reconstruction and uniqueness testing computational times with error bars of plus/minus one standard deviation over 10 repetitions at each relative sparsity and relative sampling value. Left: L1 for **signed-spikes** image class. Right: ATV for **alternating-projection** image class. Top: Size 32, middle: size 64, bottom: size 128. Legend R: reconstruction (full lines), UT: uniqueness test (dashed lines).

problems of reconstruction and uniqueness testing. The use of an interior-point method is what causes the computational time to increase so dramatically from the order of 10^0 seconds at $N_{\text{side}} = 32$ to 10^1 seconds at $N_{\text{side}} = 64$ and 10^3 seconds at $N_{\text{side}} = 128$. With another optimization algorithm shorter running times may be observed with a different result of the comparison. However, our intention with the present study is not an exhaustive algorithm comparison, but merely to demonstrate that uniqueness testing can be accomplished in the same time, or faster, than reconstruction.

7 Conclusion

The present work was motivated by understanding quantitatively how much undersampling is admitted for sparsity-exploiting reconstruction methods for CT given the lack of theoretical guarantees from compressed

sensing. Our results demonstrate empirically that sharp average-case phase transitions from no recovery to full recovery as seen in compressed sensing also occur for CT measurements across a range of image classes and sparse reconstruction methods. The location of the phase transition, i.e., the level of sampling sufficient for recovery depends on the reconstruction method and is to a large degree is governed by the image sparsity, quite independent of the particular image class.

Due to the inherently empirical nature of our study design it is clear that our results do not imply any theoretical guarantee. Further, being average-case results leaves the chance for single instance to require more or fewer samples for recovery than predicted by the average case. Nevertheless, we think the results may be used or extended to serve as guide lines for how to many CT samples to acquire based on prior knowledge about the image class and sparsity. Natural future work would be extensions toward more realistic scenarios including noisy data, model inconsistencies, specialized image classes, etc.

Constructing phase diagrams by reconstruction cannot establish solution uniqueness, which makes the uniqueness test more desirable from a theoretical perspective. However, we observed almost exact agreement between reconstruction and uniqueness test phase diagrams, so in practice the advantage may be negligible. Also, the reconstruction approach has the advantage that it can be run directly on any reconstruction problem with no need to derive specific uniqueness conditions and as such is more easily generalizable.

In our view, the presented empirical evidence suggests that an underlying theoretical explanation of phase transition behavior in CT may exist. Establishing such theory would have large implications for the understanding of sparse reconstruction in CT and we hope that the present results can serve as a step towards this goal.

Acknowledgements

The authors are grateful to Emil Y. Sidky for inspiring discussions. This work was supported in part by Advanced Grant 291405 ‘HD-Tomo’ from the European Research Council and by grant 274-07-0065 ‘CSI: Computational Science in Imaging’ from the Danish Research Council for Technology and Production Sciences. JSJ acknowledges support from The Danish Ministry of Science, Innovation and Higher Education’s Elite Research Scholarship. CK and DAL acknowledge support from the Deutsche Forschungsgemeinschaft (DFG) within the project ‘Sparse Exact and Approximate Recovery’ under grant LO 1436/3-1.

References

- [1] E. Candès. Compressive Sampling. In Marta Sanz-Solé, Javier Soria, Juan Luis Varona, and Joan Verdera, editors, *Proceedings of the International Congress of Mathematicians*, volume 3, pages 1433–1452, Madrid, Spain, 2006.
- [2] D. L. Donoho. Compressed sensing. *IEEE Trans Inf Theory*, 52:1289–1306, 2006.
- [3] E. Y. Sidky, C.-M. Kao, and X. Pan. Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT. *J Xray Sci Technol*, 14:119–139, 2006.
- [4] E. Y. Sidky and X. Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Phys Med Biol*, 53:4777–4807, 2008.
- [5] L. Ritschl, F. Bergner, C. Fleischmann, and M. Kachelrieß. Improved total variation-based CT image reconstruction applied to clinical data. *Phys Med Biol*, 56:1545–1561, 2011.
- [6] E. J. Candes and T. Tao. Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans Inf Theory*, 52:5406–5425, 2006.
- [7] D. Needell and R. Ward. Near-optimal compressed sensing guarantees for total variation minimization. *IEEE Trans. Image Process.*, 22:3941–3949, 2013.
- [8] L. I. Rudin, S Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

- [9] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Springer, New York, NY, 2013.
- [10] M. Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, New York, NY, 2010.
- [11] S. Petra and C. Schnörr. Average case recovery analysis of tomographic compressive sensing. *Linear Algebra Appl*, 441:168–198, 2014.
- [12] A. Denitui, S. Petra, and C. Schnörr. Phase transitions and cospase tomographic recovery of compound solid bodies from few projections. 2013.
- [13] J. S. Jørgensen, E. Y. Sidky, P. C. Hansen, and X. Pan. Empirical average-case relation between undersampling and sparsity in x-ray CT. 2014.
- [14] J. S. Jørgensen, E. Y. Sidky, and X. Pan. Quantifying admissible undersampling for sparsity-exploiting iterative image reconstruction in x-ray CT. *IEEE Trans Med Imaging*, 32:460–473, 2013.
- [15] D. Donoho and J. Tanner. Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philos Trans R Soc Lond Ser A*, 367:4273–4293, 2009.
- [16] T. M. Buzug. *Computed Tomography. From Photon Statistics to Modern Cone-Beam CT*. Springer, Berlin, Germany, 2008.
- [17] P. C. Hansen and M. Saxild-Hansen. AIR Tools – A MATLAB package of algebraic iterative reconstruction methods. *J Comput Appl Math*, 236:2167–2178, 2012.
- [18] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J Sci Comput*, 20:33–61, 1998.
- [19] Junzhou Huang and Tong Zhang. The benefit of group sparsity. *Ann Stat*, 38:1978–2004, 2010.
- [20] M. D. Plumbley. On polar polytopes and the recovery of sparse representations. *IEEE Trans Inf Theory*, 53:3188–3195, 2007.
- [21] M. Grasmair, M. Haltmeier, and O. Scherzer. Necessary and sufficient conditions for linear convergence of L1-regularization. *Comm Pure Appl Math*, 64:161–182, 2011.
- [22] M. Haltmeier. Stable signal reconstruction via l1-minimization in redundant, non-tight frames. *IEEE Trans Signal Process*, 61:420–426, 2013.
- [23] MOSEK ApS. MOSEK Optimization Software, version 6.0.0.122, 2011.
- [24] P. Tseng. Applications of a splitting algorithm to decomposition in convex programming and variational inequalities. *SIAM J Control Optim*, 29:119–138, 1991.
- [25] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.
- [26] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *J Math Imaging Vis*, 40:120–145, 2011.
- [27] Ernie Esser, Xiaoqun Zhang, and Tony F. Chan. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM J Imaging Sci*, 3:1015–1046, 2010.
- [28] Jeffrey D. Blanchard, Coralia Cartis, and Jared Tanner. Compressed sensing: how sharp is the restricted isometry property? *SIAM Rev*, 53:105–125, 2011.
- [29] S. Nam, M. E. Davies, M. Elad, and R. Gribonval. The cospase analysis model and algorithms. *Appl Comput Harmon Anal*, 34:30–56, 2013.
- [30] Robert Hesse and D. Russell Luke. Nonconvex notions of regularity and convergence of fundamental algorithms for feasibility problems. *SIAM J Optim*, 23:2397–2419, 2013.